aws marketplace

InfluxDB

Reviews, tips, and advice from real users






Powered by PeerSpot

Contents

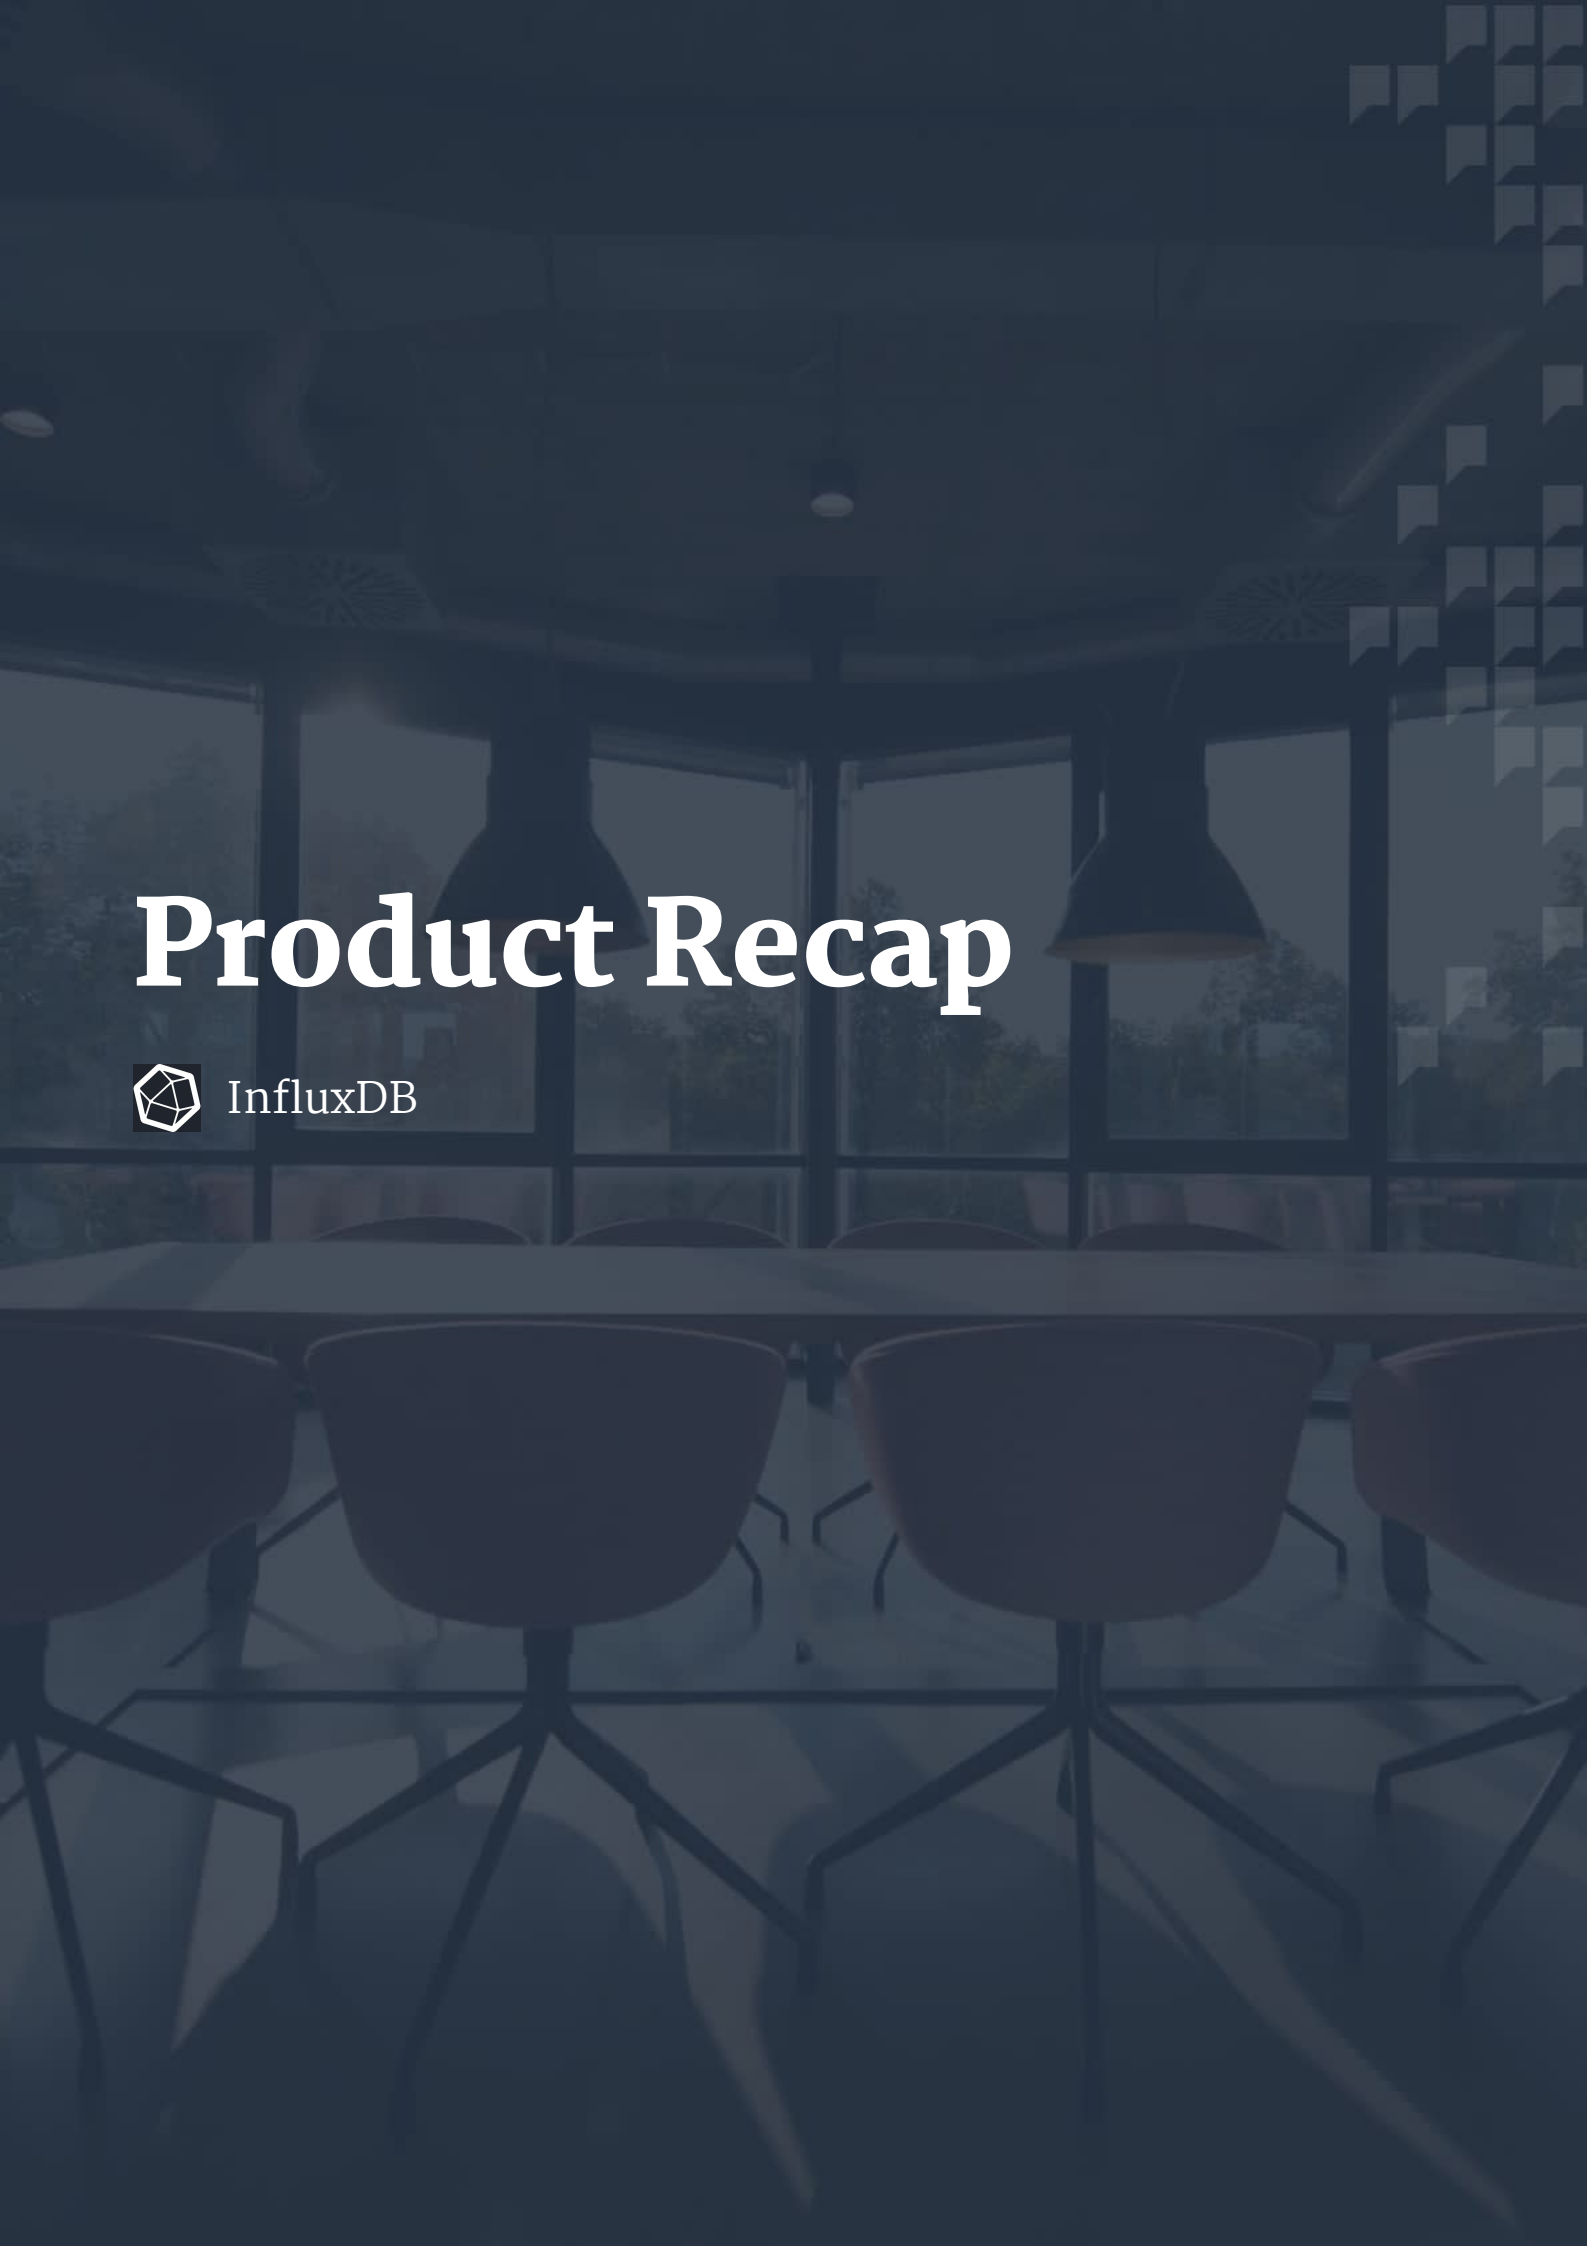

Product Recap

InfluxDB

InfluxDB Recap

InfluxDB is open-source software that helps developers and enterprises alike to collect, store, process, and visualize time series data and to build next-generation applications. InfluxDB provides monitoring and insight on IoT, application, system, container, and infrastructure quickly and easily without complexities or compromises in scale, speed, or productivity.

InfluxDB has become a popular insight system for unified metrics and events enabling the most demanding SLAs. InfluxDB is used in just about every type of industry across a wide range of use cases, including network monitoring, IoT monitoring, industrial IoT, and infrastructure and application monitoring.

InfluxDB offers its users:

- **Infrastructure and application monitoring:** Collect, process, and analyze real-time data from edge devices to help optimize distributed infrastructure.
- **IoT monitoring and analytics:** InfluxDB is designed to store large volumes of time series data and quickly perform real-time analysis on that data. Gain insights from all the sensor data and use the collected data to create and perform automated tasks.
- **Network monitoring:** Manage responsive and high-performing networks with widely distributed resources.

InfluxDB Benefits

There are several benefits to using InfluxDB . Some of the biggest advantages the solution offers include:

- **APIs and ready toolset:** InfluxDB can be accessed via a set of powerful tools enabling users to get started quickly, with less programming required. This includes a REST API, extensive client libraries, a wide variety of open-source integrations, and Flux - a functional data scripting language for querying, analysis, and events. The InfluxDB API can be used to write data from edge devices to the InfluxDB instance
- **Time series engine:** Get any data - events, logs, traces - from any edge device - systems, sensors, queues, databases, and networks. This data is stored in a powerful and high-performing engine capable of ingesting millions of data points per second.
- **Community:** InfluxDB has a large community of cloud and open-source developers ready to assist users.
- **Ready-made templates:** Use InfluxDB Templates, a set of tools with a packager and other ready monitoring solutions. These tools allow users to share their monitoring expertise with coworkers and other community members around the world. The

Templates gallery offers available templates for some of the most popular tools and applications.

- **Enhanced UI:** InfluxDB's UI includes an explorer, dashboarding tools, and a script editor. Use it to easily browse the collected metric and event data and apply common transformations. The dashboarding tool comes with a variety of visualization options that help users view insights from the data. The script editor assists users to quickly master Flux with easily accessible examples, auto-completion, and real-time syntax checking.

Reviews from Real Users

InfluxDB stands out among its competitors for a number of reasons. Two major ones are its flexible integration options and its data aggregation feature.

Shalauddin Ahamad S., a software engineer at a tech services company, notes, "The most valuable features are aggregating the data and the integration with Grafana for monitoring."

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- "InfluxDB works as expected with excellent scalability and stability, which is critical for our application."

Henning Jansen
Senior System Developer at Norled

- "While I would rate InfluxDB a ten on a scale of one to ten, users should be thoughtful about matching the engine to their specific needs."

Milan Predic
Chief Technology Officer at Presta Agency

- "InfluxDB is a database where you can insert data. However, it would be best if you had different components for alerting, data sending, and visualization. You need to install tools to collect data from servers. It must be installed on Windows or Linux servers. During installation, ensure that the configuration file is correct to prevent issues. Once data is collected, it can be sent to InfluxDB. For visualization, you can use open-source tools like Grafana."

DeepakR
Site Reliability Engineer(Observability) at Sapiens

"The user interface is well-designed and easy to use. It provides a clear overview of the data, making it simple to understand the information at hand."

BhupendraSharma

Cloud and DevOps Engineer at Betaque

"In our case, it started with a necessity to fill the gap that we had in monitoring. We had very reactive monitoring without trend analysis and without some advanced features. We were able to implement them by using a time series database. We are able to have all the data from applications, logs, and systems, and we can use a simple query language to correlate all the data and make things happen, especially with monitoring. We could more proactively monitor our systems and our players' trends."

PedroCampos

Senior DevOps at Light And Wonder

"The platform operates very quickly. It is easy to configure, connect, and query and integrates seamlessly with Grafana."

Samael Lopez

Senior Software Developer at Sherwin-Williams de Centroamérica

"The most valuable feature of the solution is we can use InfluxDB to integrate with and plug into any other tools."

PankajSingh4

Senior Specialist at Qualitest

What users had to say about valuable features:

The most important feature for us is low latency, which is crucial in building a high-performance engine for day trading. InfluxDB can handle around ten thousand messages per second, which is essential for our requirements. The solution's ability to store time series data is also significant in our crypto trading use case where time series data about prices is critical.

Milan Predic

Chief Technology Officer at Presta Agency

Read full review

"We mainly write and read data from InfluxDB. We perform very simple queries to do time series on a key, which is a unique ID of the vessel. We will select a vessel and select from time to time stamp. That's what we do. InfluxDB's core functionality is crucial as it allows us to store our data and execute queries with excellent response times.."

Henning Jansen

Senior System Developer at Norled

Read full review

“The most valuable features are aggregating the data and the integration with Grafana for monitoring.

We have been looking at the continuous query feature and think that it might solve one of our use cases, but we have not implemented it yet. We are trying to do that now..”

Shalauddin Ahamad Shuza
Software Engineer at a tech services company with 5,001-10,000 employees

Read full review

“I use Kapacitor a lot. It is the feature that allows us to create batch jobs and interact with the data that we have. That's the most interesting feature.

The query language is very similar to SQL, which is very useful because most of us have a background in SQL language.

They have the Telegraf agent that gets the data and puts it in Influx. We can also put data into Influx through the API. We are doing a lot of custom-based agents and a lot of scripts that are adding data to Influx. As a time series database, it is very powerful and lightweight. It can deal with heavy workloads very easily. That, for me, is the main advantage of Influx as compared to other databases. .”

PedroCampos
Senior DevOps at Light And Wonder

Read full review

Other Solutions Considered

"I went straight to Influx. I came across Influx, and I tried it, tested it, and learned a little bit more about it. I saw that this would do what we needed and meet our requirements.."

PedroCampos
Senior DevOps at Light And Wonder

Read full review

Previously, we used CassandraDB and ScyllaDB, a fork of CassandraDB. While these were performant, they did not store data in the time series format essential for our needs. Once we discovered that there were databases like InfluxDB designed for time series data, we decided to try it.

Milan Predic
Chief Technology Officer at Presta Agency

Read full review

"We mainly went from Icinga. It is like Nagios. So, we were using Nagios monitoring, but because the monitoring was so reactive, we wanted something new that could help us with the business requirements. We needed something that could host the metrics and data from multiple sources and can be used to correlate. It was because of the necessity we moved away from Nagios, which is very focused on systems monitoring, to something that uses a time series database and has a simple SQL query language that we can use to correlate data and create trends.."

PedroCampos
Senior DevOps at Light And Wonder

Read full review

Use Case

“I use the solution to store and manage data from various sensors in a production environment. I have developed a system where data from these sensors is communicated through an OPC UA receiver and stored in InfluxDB. It handles serial data from sensors effectively and integrates well with Grafana for visualization..”

Samael Lopez
Senior Software Developer at Sherwin-Williams de Centroamérica

Read full review

“Initially, we used Prometheus and Grafana for alerts. But, it wasn't feasible to handle logs and other types of databases. We then switched to InfluxDB due to its user interface, which allowed us to customize alerts, data visualization, and filtering for monitoring purposes. .”

BhupendraSharma
Cloud and DevOps Engineer at Betaque

Read full review

"InfluxDB is a database where you can insert data. However, it would be best if you had different components for alerting, data sending, and visualization. You need to install tools to collect data from servers. It must be installed on Windows or Linux servers. During installation, ensure that the configuration file is correct to prevent issues. Once data is collected, it can be sent to InfluxDB. For visualization, you can use open-source tools like Grafana.."

DeepakR
Site Reliability Engineer(Observability) at Sapiens

Read full review

We are developing a trading agent that uses multiple machine learning models to adapt to the crypto market in real time. InfluxDB is used to collect data on crypto coin prices from exchanges like Binance and Bybit. Our use case requires low latency and the ability to query data effectively. We use InfluxDB on a DigitalOcean infrastructure in a containerized environment with Docker.

Milan Predic
Chief Technology Officer at Presta Agency

Read full review

"InfluxDB is the main component in our large enterprise-scale streaming data application for maritime vessels. We collect position data from vessels around the coast once per second, put it on a Kafka stream, and feed those positions into InfluxDB continuously. This has been working flawlessly since 2018. We have seven years of time-series data for all the vessels that my company operates, roughly 130 to 140 vessels. Every move they make is being tracked and stored in InfluxDB.."

Henning Jansen
Senior System Developer at Norled

Read full review

"Our usage of Influx is very specific. We use it for our casino stack and applications. We use it for monitoring, logging, and metrics and trend analysis for player wagering. It is for casino data. We correlate all that data and generate reports for data warehousing. So, we use it massively in our stack for the casino for a lot of things.

We haven't yet migrated to Influx 2.0. We are still using Influx 1.8 because we have multiple workloads in production. It is going to take some time, and we need to ensure that the latest one is stable and the framework also brings a lot of different languages.

We use it on-premise and in the cloud. We have both. It is a private cloud on AWS.."

PedroCampos
Senior DevOps at Light And Wonder

Read full review

Setup

The setup process involves configuring and preparing the product or service for use, which may include tasks such as installation, account creation, initial configuration, and troubleshooting any issues that may arise. Below you can find real user quotes about the setup process.

The initial setup was straightforward, as we used Kubernetes to deploy InfluxDB. Although DigitalOcean does not offer a managed database service for InfluxDB, setting up our own container was an easy process.

Milan Predic

Chief Technology Officer at Presta Agency

Read full review

"The initial setup can be intimidating for newcomers, and there is a certain threshold needed due to the performance we get. However, once familiar with the setup, it becomes streamlined.."

Henning Jansen

Senior System Developer at Norled

Read full review

"The initial setup is straightforward. As a first-time user, the deployment usually takes us half an hour to understand all the requirements. After that, we can install the solution in one to two hours, which is efficient for installing on the client side.."

PankajSingh4

Senior Specialist at Qualitest

Read full review

"The initial setup is very straightforward. It is not very tricky.

I have tried both versions for on-premises and cloud deployment. Otis took around five to six minutes. If we have a script, cloud deployment takes one to two minutes.."

BhupendraSharma

Cloud and DevOps Engineer at Betaque

Read full review

"The initial setup was a little bit complex for a cluster. When trying to join the cluster, we were failing. We looked at the log but it was not very informative. We agreed to try a few things and after a while, we were able to join the cluster. That is the difficulty that we had.."

Shalauddin Ahamad Shuza

Read full review

Software Engineer at a tech services company with 5,001-10,000 employees

"We implemented everything. The setup is straightforward. There is a learning curve, but the setup is very easy. They have repositories for most distributions. You just have to install from a binary, and the entire stack of Influx is up and running. For most operating systems, they already have repositories that allow you to just install it with one command. I would rate it a ten out of ten in terms of ease of setup.

In terms of duration, deploying InfluxDB took minutes. It is like installing an app on your laptop. It is very easy and quick to deploy InfluxDB, but deploying what we have has been an ongoing process for the last five years.."

PedroCampos

Read full review

Senior DevOps at Light And Wonder

Customer Service and Support

We have not needed to contact technical support. All resources required were available through documentation, enabling us to resolve any issues on our own.

Milan Predic
Chief Technology Officer at Presta Agency

Read full review

"We have not contacted technical support for InfluxDB yet. When we have a problem, we consult with our other teams, and most of the time, they solve our problems.."

Shalauddin Ahamad Shuza
Software Engineer at a tech services company with 5,001-10,000 employees

Read full review

"I've never contacted technical support. If I have questions, I look at forums to find the answers that I need. I've found their forums to be very good and quite helpful.."

Bogdan Tsegelnik
Engineer at IT Specialist LLC

Read full review

“I did a search and used the community forums and help, but I never had to contact them. Their community is really big, helpful, and active. It is very easy to just go to their support site and search for solutions or ask questions. That's another thing where they are very good..”

PedroCampos
Senior DevOps at Light And Wonder

Read full review

“We had a user who worked with the support team, and they helped a lot, but they couldn't retrieve the complete data quickly.

The support team is good, but if they don't have the data, it's not their fault. It's the technical team's issue. .”

BhupendraSharma
Cloud and DevOps Engineer at Betaque

Read full review

"If you deviate from their documentation, they may refuse to provide support, stating it's not covered under the agreement. Additionally, their response times are slow, and they often suggest purchasing premium support for quicker resolution.

When you opt for premium support, they usually assign dedicated consultants. This means that whenever you encounter a problem, you have direct access to experts whom you can call, email, or engage in a call to resolve issues.."

DeepakR
Site Reliability Engineer(Observability) at Sapiens

Read full review

Other Advice

My advice for new users would be to ensure you are choosing the right engine for your domain. For InfluxDB, it performs well for low latency inputs and high-performance real-time data. While I would rate InfluxDB a ten on a scale of one to ten, users should be thoughtful about matching the engine to their specific needs.

Milan Predic
Chief Technology Officer at Presta Agency

Read full review

"Overall, I would rate the solution a seven out of ten because InfluxDB needs to remove the limitations so users can see all their data. This was a huge headache for me because we got hit twice or thrice.."

BhupendraSharma
Cloud and DevOps Engineer at Betaque

Read full review

"I recommend going with the cloud version of InfluxDB, as it's very cheap and effective. InfluxDB is very powerful, but it requires some technical skill and can't be used by just anyone. I would rate InfluxDB seven out of ten.."

Vamsi Krushna Lingala
Lead Data Engineer at Chingari

Read full review

"It depends a lot on what you are going to use it for but just enjoy it. It is very powerful. I would advise learning the TICKscript language. I know that the latest Influx 2.0 is using a different language than Influx SQL, but I would advise learning how to develop using their framework. You will see how powerful it is to use for a big data cluster. So, explore the batch processing and the scripting language that they have in 2.0. We use TICKscript, which is very powerful and allows you to do everything you need.

I would rate it an eight out of ten.."

PedroCampos
Senior DevOps at Light And Wonder

Read full review

"I give the solution a seven out of ten. The monitor is easy to use and integrate. Integration with the tools is very easy, which is why it is comfortable to test and monitor the tools that are connected to InfluxDB. we can then see the results of our monitoring.

We have over 50 people in our organization using InfluxDB and over 500 people testing.

We are using InfluxDB because it is easy to use and in my time, it has helped in live monitoring of the log testing.

For first-time users of the solution, it can be difficult to learn how to write custom data and classify it for monitoring purposes. There are many builder models in the market that can be implemented and used, but if we want to write our own, using InfluxDB can be difficult.."

PankajSingh4
Senior Specialist at Qualitest

Read full review

“If you're considering using InfluxDB for the first time, I recommend trying it. It's an open-source database with the option to purchase enterprise support. When purchasing the TICK stack, I recommend you opt for premium support. Premium support can be beneficial, particularly when encountering configuration issues or other challenges. With premium support, you can quickly reach out for assistance via phone and work towards resolving any issues promptly. Otherwise, you may wait for up to another week to resolve the problem.

Learning to use InfluxDB is not necessarily easy for a beginner. It requires some understanding, especially in areas like automation and scripting.

I would rate it as an open-source tool around seven to eight out of ten. You only need to spend money on the resources, not much on the product itself. Essentially, you need to invest in the database..”

DeepakR
Site Reliability Engineer(Observability) at Sapiens

Read full review

Top Industries

by visitors reading reviews

Financial Services Firm	12%
Computer Software Company	11%
Comms Service Provider	9%
Manufacturing Company	9%

Company Size

by reviewers

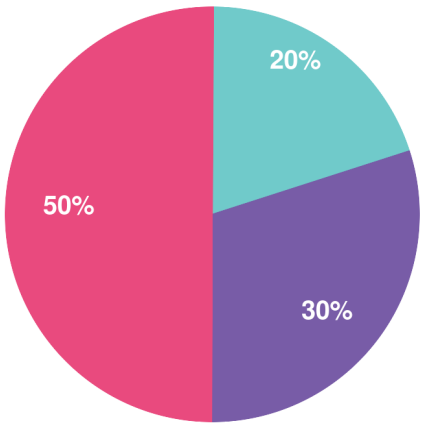

by visitors reading reviews

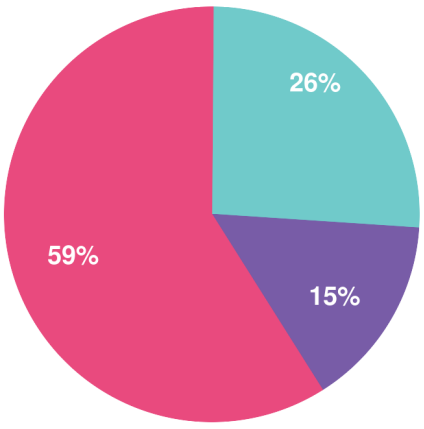

Large Enterprise Midsize Enterprise Small Business

About this buyer's guide

Thanks for downloading this PeerSpot report.

The summaries, overviews and recaps in this report are all based on real user feedback and reviews collected by PeerSpot's team. Every reviewer on PeerSpot has been authenticated with our triple authentication process. This is done to ensure that every review provided is an unbiased review from a real user.

Get a custom version of this report... Personalized for you!

Please note that this is a generic report based on reviews and opinions from the collective PeerSpot community. We offer a customized report of solutions recommended for you based on:

- Your industry
- Company size
- Which solutions you're already considering

The customized report will include recommendations for you based on what other people like you are using and researching.

Answer a few questions in our short wizard to get your customized report.
Get your personalized report here

About PeerSpot

PeerSpot is the leading review site for software running on AWS and other platforms. We created PeerSpot to provide a trusted platform to share information about software, applications, and services. Since 2012, over 22 million people have used PeerSpot to choose the right software for their business.

PeerSpot helps tech professionals by providing:

- A list of products recommended by real users
- In-depth reviews, including pros and cons
- Specific information to help you choose the best vendor for your needs

Use PeerSpot to:

- Read and post reviews of products
- Access over 30,000 buyer's guides and comparison reports
- Request or share information about functionality, quality, and pricing

Join PeerSpot to connect with peers to help you:

- Get immediate answers to questions
- Validate vendor claims
- Exchange tips for getting the best deals with vendor

Visit PeerSpot: www.peerspot.com

PeerSpot

244 5th Avenue, Suite R-230 • New York, NY 10001

reports@peerspot.com

+1 646.328.1944